

Development and Implementation of Large-Scale Micro-Robotic Forces Using Formation Behaviors

D. D. Dudenhofer*, D. J. Bruemmer**, M. O. Anderson***, and M. D. McKay****
Idaho National Engineering and Environmental Laboratory (INEEL)
P.O. Box 1625, Idaho Falls, ID 83415

ABSTRACT

Micro-robots may soon be available for deployment by the thousands. Consequently, controlling and coordinating a force this large to accomplish a prescribed task is of great interest. This paper describes a flexible architecture for deploying thousands of autonomous robots simultaneously. The robots' behavior is based on a subsumption architecture in which individual behaviors are prioritized with respect to all others. The primary behavior explored in this paper is group formation behavior drawn from the work in social potential fields applications conducted by Reif and Wang¹, and Dudenhofer and Jones.² While many papers have examined the application of social potential fields in a simulation environment, this paper describes the implementation of this behavior in a collective of small robots.

Keywords: Robot, Social potential field, Online learning, Swarm behavior, Distributed Robotics, DOE, RIM

1. INTRODUCTION

Technological advances in micro-robotics, remote sensors, and artificial intelligence continue to increase the capabilities of micro-robots while decreasing the size of such units. Within the near future, it may be possible to produce and deploy thousands of inexpensive, essentially disposable micro-robots. Although possibly limited in individual capability, robots deployed in large numbers represent a tremendous cumulative force. Given the proper social behavior set, the agents form a collective much like a colony of ants or swarm of bees. Importance shifts from the actions of individual agents to the collective behavior. Therefore the key to success is developing the individual robot behavior sets which promote the emergence of desired collective behaviors. One such desired behavior is that of coordinated group motion within a structured formation. A structured or group formation in this context does not necessarily refer to adherence to a strict geometric pattern such as a circle, square, diamond, etc. Rather, as observed in nature, group formation behavior promotes a spatial relationship between adjacent entities such as might be seen in a flock of birds, a school of fish, or even a swarm of ants. In all of these cases, the entity, be it a fish, bird or insect, does not possess a global knowledge, but rather senses and reacts to its immediate neighbor. From this simple concept arises many complex emergent behaviors. This paper examines Social Potential Fields as a means to promote and control a variety of these emergent effects in a team or swarm of small robots. Our work with a collective of 12 small robots shows that group formation wrought entirely through local interactions and reactive behaviors can provide a means for global coordination and control of a collective as it performs searches in various environments. Through simple behavior adjustments made either through online adaptation or in response to high-level user commands, it is possible to spur dramatic changes in the behavior of the collective. We have also shown that the ability to appropriately modulate these changes can provide significant improvements in the performance of the collective. In addition to examining the utility of Social Potential Forces, this paper also discusses issues of integrating such a system in hardware.

1.1 Relevance To Doe Objectives

In *Robotics and Intelligent Machines in the U.S. Department of Energy: A Critical Technology Roadmap*,³ the Department of Energy (DOE) has identified the potential for distributed robots to play a role in reducing cost, improving worker health and safety, augmenting product quality and increasing overall productivity. Robots will be constructed to cost-effectively perform complex, hazardous tasks that humans cannot or should not have to do. We envision the robotics technologies developed in our research being used to map and characterize buried waste and retired facilities; to perform routine inspection and monitoring of critical components; to perform environmental monitoring and building surveillance appropriate for DOE long term stewardship needs and to provide rapid response remote characterization capabilities in the event of a hazardous spill or radiation leak. Small scale distributed robot systems, such as the one discussed in this paper, can reduce cost, remove workers

Correspondence: * dudedd@inel.gov, phone: 208-526-0700, fax: 208-526-0777, ** bruedj@inel.gov, phone: 208-526-4078, *** matthew@inel.gov, phone: 208-526-4308, **** mark@inel.gov, 208-526-0539, Idaho National Engineering and Environmental Laboratory (INEEL), P.O. Box 1625, Idaho Falls, ID 83415. This work is supported through the INEEL Long-Term Research Initiative Program under DOE Idaho Operations Office Contract DE-AC07-99ID13727 and through DARPA Software For Distributed Robotic program contract J933.

from the dangers of radioactive or hazardous materials, and increase productivity by accomplishing slow painstaking tasks. Paramount to realizing these benefits is the construction of robust robot behaviors coupled with human operator interface systems, which promote system understanding and facilitate human interaction.

1.2 Project Objective

The motivation of this project is to develop a team of small disposable robots to assist a human operator in the remote characterization of hazardous or unknown environments. Numerous scenarios support the need for this type of capability. Examples include characterization and assessment following an explosion at a hazardous material processing plant or a radioactive spill at a nuclear power facility. Similarly, this technology could be employed in an urban warfare setting in conducting a preliminary search for hidden weapons of mass destruction, chem/bio hazards or terrorists in a suspect building.

The goal of using autonomous robots to perform remote characterizations is not unique. Current robotic systems used for this application, tend to be highly sophisticated, expensive platforms – typically large to mid-sized robots deployed as single units or in small groups. Existing systems are domain-centric and often require prior instrumentation and/or teams of specialists to operate and service. They are expensive to manufacture, transport, and operate, and, consequently, are undesirable for rapid response in remote characterization tasks where the robots often cannot be recovered because of exposure to hazards. One innovative approach to the remote characterization problem was developed through a partnership between Oak Ridge National Laboratory (ORNL) as with the Idaho National Engineering and Environmental Laboratory's (INEEL). This system employs MACS (Mobile Automated Characterization System) as a mid-sized robot which explores and characterizes the building, deploying a smaller robot, RACS (Reduce Access Characterization System), into rooms and areas where there is limited mobility/access.⁴

While the MACS and RACS team offers many benefits over strategies that employ only one system, such deployer approaches do not exploit the benefits of fully distributed systems. Specific benefits of distributed systems with a large number of smaller robots include:

- *Emergent Behavior* – As in a colony of ants, intelligent, complex behavior emerges from the interactions of multiple robots each driven by simple behaviors.
- *High Fault Tolerance* – By distributing the task across a loosely coupled population of robots, the collective can succeed even when particular robots fail.
- *Redundancy* – The behavior of each robot can be validated / duplicated by its peers.
- *Cooperative Behavior* – We can exploit synergistic behavior impossible with only one to several robots.
- *Modulated Diversity* – As in biological systems, an appropriate level of diversity adds richness to the capabilities of the collective and makes it more robust to environmental changes.
- *Low Cost* – Small scale robots can potentially be used as a disposable resource

These characteristics are desirable in promoting distributed robotic systems as a viable and preferred method for remote characterization over current methods which use expensive single unit systems or use direct human control to conduct the characterization.

1.3 Platform Selection

Desiring a robust deployment system with the characteristics as described above, this project centers on the deployment of teams of small to micro-sized robots as an application platform for conducting survey and assessment missions. This paper does not suppose that large forces of small robots are the panacea for all missions, but their employment has definite strengths for use in large area coverage and / or hazardous / single use environments. We believe our approach is especially useful for rapid response type missions where there is little known about the environment or insufficient time to custom tailor a robotic solution. The subset of robots and robot systems under evaluation possess the following individual and collective characteristics:

- Small in size with a cross section between 2 to 4 inches.
- Cheap with ultimate production costs less than \$100 per unit.
- Diminutive in individual capability, but collectively powerful in information collection.
- Reactive in nature possessing only a local awareness.
- Autonomous – possessing a self-sufficient behavior set.

- Interactive – possessing the ability to exchange information with other robots and / or human operators.
- Able to be commanded by human operators who may modify the robot’s behavior to supplement local knowledge, as more intelligence becomes available or tasking changes.
- Deployed in medium to large groups.
- Capable of interfacing to various mission-specific sensors such as moisture or radiation sensors.
- Non-GPS reliant. Although GPS or a similar positioning system may ultimately be fitted to the platform, no reliance on an absolute, centralized positioning capability is required.
- Limited mobility (i.e. system is designed for distances under 100 meters, smooth terrain)
- Limited in individual communications range.

The test platform chosen for this research project is the Growbot by Parallax. It is an “off the shelf” educational robot kit that typically sells for \$175 to \$250 per unit depending on the configuration. Compared to larger robots, the Growbot is extremely limited in capability. This was a specific research choice and a research challenge, to create a powerful collective of individually diminutive robots. Figure 1 below shows one of our robots in a basic configuration with a moisture detection sensor, bump sensor, two whisker-like light sensors, and four IR sensors for obstacle avoidance. A subset of our robot team is also configured with IR transmission and RF transmission capabilities, and a specialized speaker/microphones.

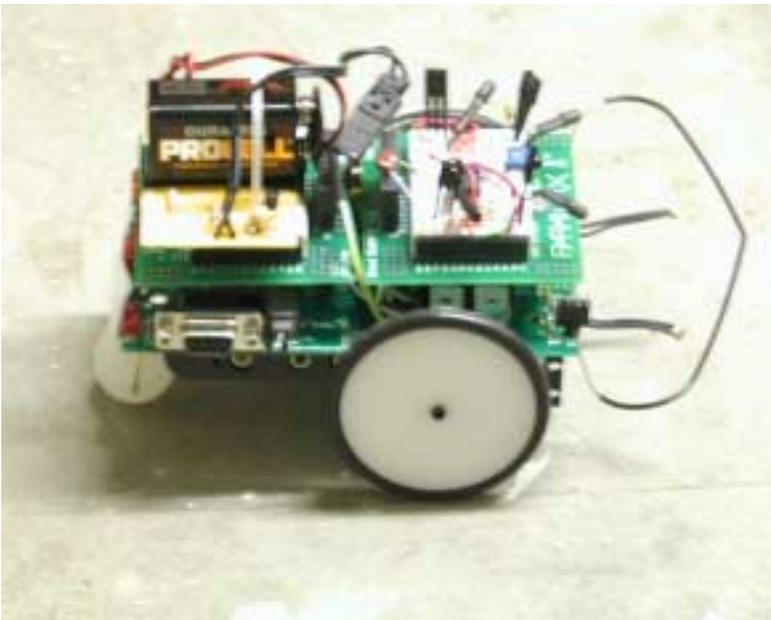


Figure 1: Growbot

Basic Stamp 2 Characteristics	
Cost:	\$50 one per robot
Processor:	microchip pic16c57
Package:	24-pin dip
Environment:	0 deg to 70 deg C (operating temp)
Clock speed:	20Mhz
Ram:	32 bytes ram
Program size:	2K EEPROM
Voltage:	5-15 VDC
Current draw:	8mA (running) 100microA (sleep)

Table 1: Basic Stamp Characteristics

Table 1 lists the characteristics of the Basic Stamp 2. It is easy to see that compared with many other processors and chip configurations, the Basic Stamp is very diminutive in capability. Within these limitations, however, we constructed a collection of robots, which are domain general and robust to environmental changes for the very reason that they are simple. The robots are designed to respond to touch, light, IR break beams, and sound. Additionally, the robots exhibit self-learning in which they constantly adjust their reliance on these sensors based on their experiences (e.g. how many bumps, ambient light, etc.). The project objective is to create a force of Growbots equipped with sensors and a behavior set which could enter a hazardous area and identify the boundaries to a hazardous material covering a concrete floor. This paper examines the research to date with specific attention to the implementation of the formation behaviors in the robots.

2. INDIVIDUAL AND COLLECTIVE BEHAVIORS

One of the key elements in developing a distributed robotic system is designing a behavior set that promotes the desired behavior in the individual robots while at the same time facilitating the desired collective behavior. Insight into this development can be gained through observing nature. Social communities in nature, i.e. bees, ants, birds, fish, provide the inspiration and insight for much of the work in collective robotics. Of particular interest is the ability that birds and fish have to form and maintain collective patterns even in the face of dynamic environments. These patterns are formed by the

animal's ability to balance the desire to remain close to the flock (or school) and also to avoid collision.⁵ Within the flock, the bird does not possess universal knowledge (i.e., knowledge of the position of all others in the flock), but it adjusts its position based on the perception of its immediate neighbors. A swarm of ants is another good example. While each individual ant's behavior may seem chaotic, the collective behavior has been fine-tuned through evolution to ensure the survival of the species. Our approach pulls from these observations in nature and in assigning each individual robot a simple set of individual behaviors. Within this framework, each robot senses and reacts to its environment and other neighboring robots. Together the robots function as a collective without the use of any behavior specifically dedicated as such.

To address the issue of individual and collective motion, we have integrated social potential based behaviors within our robots. Arkin⁶ describes social potential fields: "This method generates a field representing a navigational space based on an arbitrary potential function. The classic function used is that of Coulomb's electrostatic attraction, or analogously, the law of universal gravitation, where the potential force drops off with the square of the distance between robots and objects within its environments." The actual response of a robot is based on the cumulative force of the sensed environment. Recent examples of the application of potential fields for multi-robot formation tasking include the work by Balch and Hybinette.⁷ In their scheme, every robot possesses local attachment sites at which social potential fields promote preferential alignment by neighboring robots. This shaping mechanism has produced diamond, line, and square formations for small groups of simulated robots. The work by Schneider, Wildermuth, and Wolf⁸ discusses a method for employing social potential fields in developing specific geometric patterns, which were demonstrated in simulation.

While these two efforts have focussed on the formation of specific geometric patterns, this project focuses on the use of social potential fields as a method of dispersion and coverage control. This project builds upon the work with simulated robots conducted by Reif and Wang¹, and Dudenhoeffer and Jones,² which examined the application of social potential field on large-scale numbers of robots to promote swarm-like behavior.

3. IMPLEMENTATION IN ACTUAL ROBOTS

The hardware implementation of social potential fields in a team of Growbot robots required the development of a scheme which provided a repelling forces to promote spatial separation during an area search, and an attractive forces as a means for pulling the collective together. This implementation in hardware was not a trivial matter, especially given the limited capability of the individual robots in terms of sensors, communications and processing. The robots possessed no knowledge of self-location nor did they have the capability to identify discrete ranges to other robots in the team. This is contrary to other work in the implementations of social potential, where the robots possess GPS, sonar, lasers, and direct communications with neighboring robots. Here the challenge was not only to develop individual behaviors, but also to develop individual behaviors which promoted the desired collective behavior; i.e. an area search for hazardous materials. The development and employment strategy centered on a team of 12 robots. The robots are autonomous, react to their environment and fulfill basic search and exploration tasks. Additionally, an interface structure was included in the system to allow high level command from a human operator.

To efficiently enable this communication, the swarm was subdivided into two types of robots, Sergeants and Privates. The Sergeants are specialized robots configured to lead and reconfigure groups under their control. Since the robots do not form a communication network, the Sergeants accomplish this by displacing communication and control onto the environment in the form of ambient "chirps", distinct audible tones. The advantages of chirps are that they are low-power, unambiguous, not limited to line of sight, scale to large numbers of robots, require no communication protocols or synchronization, and by their intrinsic physical characteristics (e.g. volume, direction of source), implicitly encode position information. Chirps are used by the Sergeants to maintain dispersion gradients and guide group movement. The "chirps" represented the attraction arm of the social potential field implementation. Chirps are also used by each robot in the swarm as a "come hither" call to bring about convergence once a spill is found. Each robot in the swarm uses a subsumption-style architecture consisting of basic behaviors including obstacle avoidance, follow "command" chirp, "converge to spill" chirp, and a spill seeking behavior.

Privates are intended to operate at a level similar to simple insects in that they can autonomously avoid obstacles and search for spills, but perceive and interact only with their local environment. The Sergeants have all the characteristics of Privates, but are also outfitted with bi-directional RF communication hardware. A RF communications link allows human operators to inject intelligence and high level commands through a command and control software suite called *AgentCommand*, developed at the INEEL. The Sergeants pass on movement commands implicitly to the swarm by emitting the "follow me" chirp mentioned above. This command chirp beacon insures that the group remains clustered and that it moves in unison to

follow the Sergeant. Additionally, Sergeants can pass information to specific individual or specific groups of individuals via an infrared (IR) link. This allows human operators to work through Sergeants in reconfiguring the swarm and reallocating resources. The goal is not tele-operation, but a level of adjustable autonomy to support the injection of higher level intelligence into the collective.

The repelling arm of the social potential field is provided through the combined use of light sensors and IR sensors. These sensors also act as a method of collision avoidance for the robots. In our implementation, social potential fields are not independent of other behaviors such as obstacle avoidance. Rather, lower level behaviors based on light sensing, IR and chirp following work together to produce the desired effects sought in social potential fields. Collision avoidance was also supplemented through input from a bump sensor. With the use of real robots, new issues arise such as types of flooring, lighting conditions, and the presence, position and size of obstacles, sensor variability. The remainder of this paper focuses on the implementation and testing of the repelling force for maintaining spatial separation.

4. EXPERIMENTAL TEST-BED

One of the problems for gathering empirical data on the behavior of robots has been the difficulty and cost associated with using an absolute, accurate indoor positioning system to capture displacement for multiple robots over long periods of time. GPS is not suitable for such indoor purposes and other indoor positioning systems either require costly instrumentation (magnetic field beacon technology, radio, vision) or are vulnerable to drift (e.g. dead-reckoning). Besides, our robots were neither physically large enough nor computationally powerful enough to support sophisticated instrumentation or dead reckoning.

For our purposes, we constructed an environment consisting of an eight by eight foot walled enclosure with a floor covering consisting of large sheets of white-board. Each robot was instrumented with a Velcro sponge pad, which allowed us to securely attach a dry erase marker to the rear of the robot. Figures 2 and 3 illustrate the test-bed environment. The marker provided an excellent means to visually track the movements and area coverage for the robots as they explored the environment. When trials involved teams of robots, each was fitted with a different color marker to differentiate its path from the other robots. This novel approach provided valuable feedback on the robots' search strategy and the cumulative affect of adding additional robots to the search.

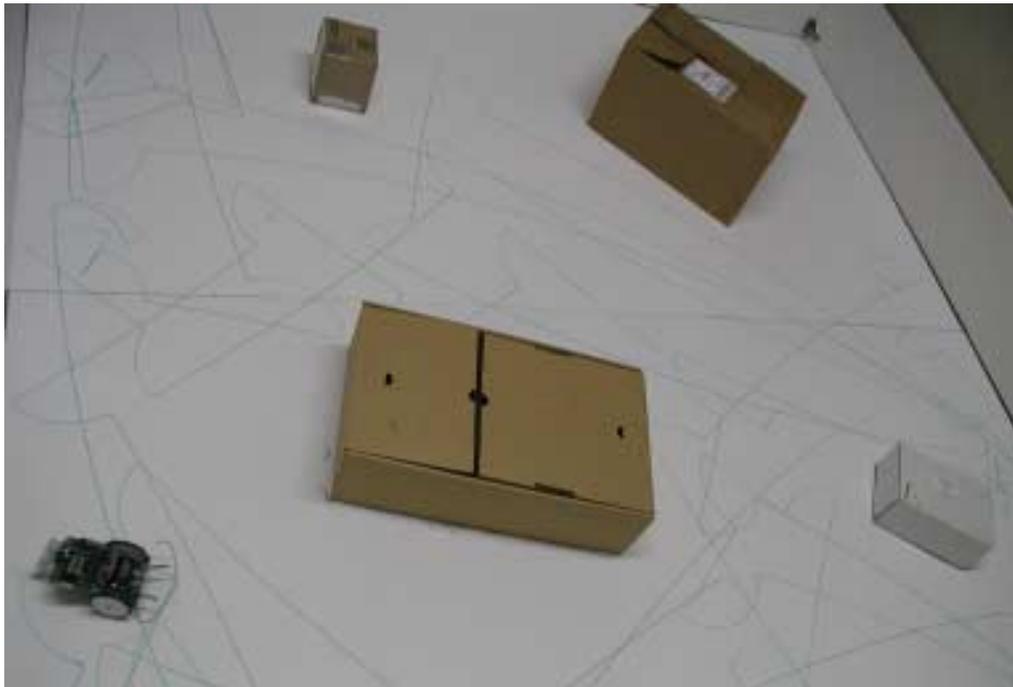


Figure 2: A single GrowBot within the testbed environment after approximately five minutes



Figure 3: A test run of a single robot after approximately 60 minutes.

With this test-bed in place, we began to address our fundamental problem of understanding nuances of collective behaviors. Increasingly, we realized that none of our hard coded search strategies were robust across different environments. For example, search strategies that worked well in tight, cluttered environments did poorly in open arenas. Likewise those strategies, which could efficiently cover wide-open areas, did poorly in tight quarters. Even behaviors which had worked well for individual robots in terms of obstacle avoidance, frontier based exploration and efficient ground coverage, proved problematic when instantiated in multiple robots. More often than not, the behavior sets we implemented produced the following undesirable emergent effects rather than the harmonic synergy we desired:

1. *Interference*:
 - Physical Interference: The robots sometimes collided or became physically entangled with one another
 - Chattering: A phenomenon whereby robots hem each other in and, given sufficient obstacle and population density, essentially spin in place. Chattering wastes time and energy and hinders exploration of new ground. This is a phenomenon characteristic of social potential forces in systems. ² It can be mitigated through omni-directional sensing, moment-type forcing function, or online learning.
2. *Redundancy*: Robots tended to cover the same ground as their peers and fall into “ruts.”
3. *Area Omission*: Each behavior tended to have its own Achilles’ heel – a set of environmental conditions which it tended to avoid such as corners, shadowed/bright areas, or areas behind obstacles.

Before we could really begin to appreciate the advantages to be had through swarming behavior, we need some way to escape the undesirable emergent effects discussed above. Our first implementation of a social potential field approach had not provided us with this means.

What we needed was a search behavior which could implicitly (i.e. without use of a map, internal representation, directed communication, or centralized control):

- Reduce redundancy and interference
- Maintain a beneficial level of social interaction
- Adjust each robot’s willingness to explore
- Automatically adapt individual robot behavior to different environments and varying numbers of robots

To accomplish these aims, we developed a form of online adaptation that could automatically adjust a triumvirate of light-sensitivity, turn gain and, at a higher level, the robot's average turn frequency. By raising these parameters the robots become more responsive to their environment, maintain a greater social potential field, and become less likely to explore new ground or escape from the swarm. By lowering these parameters the robots become less sensitive and therefore more likely to venture out, away from the swarm, and cover new ground. In the end, we arrived at a "motivation" system that we could move up and down a sliding scale of 'timid' to 'brave.'

We believe such an approach is uniquely appropriate for distributed systems especially as we scale towards large numbers of very small-scale, resource constrained robots. For one thing, the sensors involved make minimal demands in terms of cost, size and processing power. Also, the existence of such an intrinsic "motivation" system can permit a user to control swarming behavior at a high, abstract level. Most importantly, it provides the swarm with a means to automatically regulate itself. Bonabeau et al. State⁹ that artificial, swarm-intelligent systems must exhibit the features that have made social insects so successful: flexibility, robustness, decentralized control and self-organization. Although our initial system was decentralized, it lacked self-organization and flexibility and therefore also lacked robustness. In particular, Bonabeau et al. argue that self-organization must include the following four ingredients: positive feedback, negative feedback, amplification of random variations, and multiple interactions.

Whereas our initial implementation had only the multiple interactions, the online learning system now supplied the other necessary characteristics. Positive and negative feedback is supplied by a critic, invoked at given time intervals, which balances the need to explore against the need to avoid obstacles and maintain a social potential field. For instance, motivation is increased when the number of turns in a given time segment are above a threshold and lowered when the number of bumps is above a given threshold. This motivation is itself based on a form of sensor adaptation that amplifies variations in light. Essentially, this is a variation on an old theme originally developed by Grey Walter¹⁰ in 1953 and adopted in Valentino Braitenberg's¹¹ vehicles in 1984. Both Walter and Braitenberg used an inhibitory and excitatory influence based on light to produce ostensibly aggressive or cowardly behavior.

It is no coincidence that many insects use perception of light intensity to modulate behavior. Light provides continuous, readily available fluctuations that are truly random in the sense that readings fluctuate chaotically and yet meaningful in the sense that light intensity mirrors aspects of the real world such as the presence of obstacles and other agents. Perception of real world light intensity offers a perfect means to draw an appropriate level of randomness into the robots' behavior, providing a means of self-organization, and, ultimately, a means to realize the benefits of swarm intelligence.

Good search strategies find an appropriate balance of randomness and uniformity. Too much randomness and the system will be haphazard and desultory. Too little randomness and the search may perform well in some regions but exclude others, arriving at a local rather than global performance maximum. The online learning provided this needed element of randomness. By adjusting the influx of randomness, our behavior modulation system can control certain emergent properties of the swarm such as swarm density, swarm translation, and swarm convergence. It also provides a means to adapt the swarm to new environments and promotes full coverage especially in obstacle rich environments. This ability to harness chaos can help the collective avoid the problems of redundancy, interference and areas of omission. Our ability to draw an appropriate level of randomness reduces redundancy. Although all robots possess the same behavior set, online learning modifies each robot based on its unique experiences, resulting in an appropriate level of social entropy. This diversity plays a crucial role in avoiding redundancy and eliminating areas of omission. Also, as online adaptation adjusts the "timid" to "brave" motivation system, it automatically modulates the emergent swarming behavior, allowing robots to escape the pitfalls of chattering. Finally, we had a system where it seemed that robot interaction was an asset, rather than a liability.

To verify that the online adaptation was working effectively, multiple trials were evaluated with a single robot in two different domains. Environment one was replete with non-systematic, drastically non-uniform variations including bright light sources, paint splotches on the floor, and a variety of different colored and textured obstacles. Figure 4 illustrates the non-systematic environment used for evaluation. Environment two was represented by the test-bed environment and consisted of a perfectly white floor with no perceptible variation. We performed tests with one robot in both environments with the behavior modulation fixed at different points along a spectrum from "timid" to "brave." Without learning, the brave robots did well in the first environment but poorly in the second where they hit many obstacles. Conversely the timid robots did well in the second environment but remained unwilling to explore in the first.



Figure 4: The team of robots operation in the non-systematic environment.

Next learning was incorporated in which the level within the “timid” to “brave” spectrum was adjusted based upon the robot’s experience in the environment. With learning, the robots were able to adapt within two minutes to whatever environment they were placed in *regardless of what setting was chosen a priori*. Results showed that robots in environment one acquired a sensitivity level almost twice the sensitivity of those in environment one. The fitness critic used to produce this adaptation had managed to balance coverage efficiency with the robot’s ability to effectively avoid obstacles. For both environments the average total bump count for each 5-minute run was approximately two bumps.

5. COMMAND AND CONTROL ISSUES

Much progress has been made to develop control strategies for multiple mobile robots. However, for truly distributed approaches to be successfully fielded, additional research is needed to develop human interface tools that allow a user to effectively task large numbers of robots. INEEL has completed the first phase of investigative research into this area and is now embarking on a full-scale implementation effort that will bring their findings to bear on real world robots. This project incorporates the modular robot interface suite developed at the INEEL called *AgentCommand*. A user should not have to send detailed commands to individual robots; rather *AgentCommand* permits operators to interface with robot teams at a user-defined level of abstraction. The power of large-scale, distributed robot teams is not the contribution of individual robots, but rather the cumulative effect. *AgentCommand* allows the user to think and plan in terms of capabilities and functionality of groups rather than individual entities. The aim is to infuse operator intelligence into the collective at an appropriate level, in order to enable a form of adjustable autonomy that balances the strengths of both man and machine.

The implementation of Social Potential Fields combined with online learning provide a mechanism for a human operator to influence high level swarm behavior. By adjusting the “timid” to “brave” spectrum thresholds, the operator can influence the swarm size, density, and its willingness to explore. Likewise, the Sergeant / Private implementation provide a hierarchical approach for command and control. This structure eliminates the need for a human operator to directly interface with each robot on an individual level. Commands are simply passed to the Sergeants and flow down to the Privates.

6. SUMMARY AND CONCLUSION

The deployment of distributed robotic systems relies on the development of robust behaviors to promote the desired individual and collective behaviors. Spatial relation between neighboring robots is a key component in using distributed robot teams for conducting remote area survey and characterization. This project has demonstrated in real robotic platforms that social potential fields provide a viable means to promote this behavior. While this implementation of social potential fields would perhaps have little purpose on larger platforms, it can be achieved on small, very limited robots to achieve significant performance improvements. If unable to adjust to new environments and changing numbers of robots, the social potential field method can produce undesired side effects and by itself may not be sufficient to improve collective performance. The onboard learning techniques discussed in this paper provide a viable means to take advantage of the strengths of social potential fields while helping to mitigate some of its weaknesses. Automatic behavior modulation can limit detrimental effects such as chattering and physical interference. Although each robot begins with the same program, each will perform differently based on its own experiences with the environment. This diversity helps combat the problems of redundancy and area omission during searches.

In addition to studying the effects of online adaptation and behavior modulation we have begun experiments intended to provide empirical information regarding the advantages of multiple robot strategies. The development and use of a simple but effective test-bed environment provided valuable insight and understanding into collective behaviors. Such test-beds, which can provide empirical data and scale to large numbers of robots, are necessary if we are to answer key questions pertaining to multi-robot systems. What is the measurable effect of adding more robots? At what point do the benefits reaped, per robot, begin to decline? Preliminary results indicate that the use of multiple robots on an area coverage task can augment performance by a measure significantly greater than the area coverage performance of one robot multiplied by total of robots used. This data upholds the intuition that there is something significant to be reaped from the interactions between robots.

The advantages of a distributed approach may, in fact, be most significant in their ability to alter the worst case scenario and lower uncertainty. In fact, our research shows that our fully distributed approach tends to move the performance of the collective towards a more stable average case. In other words, the more robots used, the more reliable the behavior of the collective. Consequently, distributed strategies permit performance guarantees in a way that single autonomous robots cannot. For operational purposes, this may be a crucial finding. All too many promising robotics applications are stymied by the tyranny of an insidious worst case possibility. Another means to avoid the worst case and enable real-world applications is to develop interface technologies, which allow human insight to be injected at an appropriate level. For this purpose, we are continuing in the development of *AgentCommand* for high-level command and control of the collective and are currently interfacing these tools with the robot collective.

ACKNOWLEDGMENTS

This work is supported through the INEEL Long-Term Research Initiative Program under DOE Idaho Operations Office Contract DE-AC07-99ID13727 and through DARPA's Software for Distributed Robotics project under research contract J933.

REFERENCES

1. J. H. Reif and H. Wang. "Social potential fields: A distributed behavioral control for autonomous robots". *Robotics and Autonomous Systems* **27**: pp. 171-194, 2000.
2. Dudenhoefter, Donald D. and Michael P. Jones, "A Formation Behavior for Large-Scale Micro-Robot Force Deployment". In *Proceedings of the 2000 Winter Simulation Conference*, editors J. A. Joines, R. R. Barton, K. Kang, and P. A. Fishwick, PP. 972-983, Orlando FL, December 2000.
3. Department of Energy. "Robotics and Intelligent Machines in the U.S. Department of Energy: A Critical Technology Roadmap", <http://www.ncsu.edu/IEEE-RAS/RAS/RASnews/990714RIMfinal.html>
4. M. O. Anderson, M.D. McKay, and B.S. Richardson, "Multirobot Automated Indoor Floor Characterization Team". *Proceedings of the 1996 IEEE International Conference on Robotics and Automation*, pp. 1750-3, 1996.
5. E. Shaw, "Fish in schools," *Natural History* **84** (8): p. 4046, 1975.
6. Arkin, Ronald C. *Behavior-Based Robotics*. p. 73, Cambridge, Massachusetts: The MIT Press, 1998.

7. T. Balch and M. Hybinette, "Social Potentials for Scalable Multi-Robot Formations" in *Proceedings of the 2000 IEEE International Conference on Robotics & Automation*, editors C. Laugier and G. Lee, pp. 73-81, San Francisco, April 2000.
7. F.E. Schneider, D. Wildermuth, and H.L. Wolf, "Motion Coordination in formations of multiple mobile robots using a potential fields approach", *Distributed Autonomous Robotic Systems, Vol 4*, editors L.E Parker, G. Bekey,
8. and J. Barhen, pp. 305-314, Springer-Verlag, Tokyo, 2000.
9. E. Bonabeau, Dorigo, M, Theraulaz, G. *Swarm Intelligence*, pp. 7-13, New York, NY: Oxford University Press, 1999.
10. W.G. Walter, 1953. *The Living Brain*, New York, NY: Norton, 1953.
11. V. Braitenburg, *Vehicles: Experiments in Synthetic Psychology*, Cambridge, MA: MIT Press, 1984.